Grover Adaptive Search for Constrained Polynomial Binary Optimization

Austin Gilliam¹, Stefan Woerner², and Constantin Gonciulea¹

¹JPMorgan Chase

²IBM Quantum, IBM Research – Zurich

In this paper we discuss Grover Adaptive Search (GAS) for Constrained Polynomial Binary Optimization (CPBO) problems, and in particular, Quadratic Unconstrained Binary Optimization (QUBO) problems, as a special case. GAS can provide a quadratic speed-up for combinatorial optimization problems compared to brute force search. However, this requires the development of efficient oracles to represent problems and flag states that satisfy certain search criteria. In general, this can be achieved using quantum arithmetic, however, this is expensive in terms of Toffoli gates as well as required ancilla qubits, which can be prohibitive in the near-term. Within this work, we develop a way to construct efficient oracles to solve CPBO problems using GAS algorithms. We demonstrate this approach and the potential speed-up for the portfolio optimization problem, i.e. a QUBO, using simulation and experimental results obtained on real quantum hardware. However, our approach applies to higher-degree polynomial objective functions as well as constrained optimization problems.

1 Introduction

Using the laws of quantum mechanics, quantum computers offer novel solutions for resource-intensive problems. Quantum computers are theoretically proven to solve certain problems faster than a classical device [1–3] and are well-equipped to handle tasks such as factoring [2], linear systems of equations [4, 5], Monte-Carlo simulations [6–9], as well as combinatorial optimization problems [10–14].

A commonly-studied class of combinatorial optimization problems are *Quantum Unconstrained Binary Optimization* (QUBO) problems, with applications in resource allocation, finance, machine learning, and partitioning. There are multiple approaches to solve QUBO problems on a quantum computer, discussed in the following paragraphs.

First, quantum annealing [15, 16] is a meta-heuristic for adiabatic quantum computers. Similar to simulate annealing, it can be used to approximate

optimal solutions of QUBO problems.

Further, there exist variational quantum optimization heuristics, such as *Variational Quantum Eigensolver* (VQE) and *Quantum Approximate Optimization Algorithm* (QAOA) [10–13]. VQE and QAOA are heuristics designed for near-term, noisy quantum computers without performance guarantees. However, for QAOA, it is known that in the infinite depth limit, the algorithm recovers adiabatic evolution and would converge to the optimal solution.

Last, there are Grover-based [1] optimization algorithms, such as *Grover Adaptive Search* (GAS) [17–19]. GAS iteratively applies Grover Search to find the optimum value of an objective function, by using the best-known value as a threshold to flag all values smaller than the threshold in order to find a better solution. The algorithmic framework comes with a quadratic speed-up, however it likely requires an error-corrected fault-tolerant quantum computer due to the depth of the resulting circuits. One of the challenges inherent in GAS is the creation of efficient oracles.

In this paper, we provide a framework for automatically generating efficient oracles for solving *Constrained Polynomial Binary Optimization* (CPBO)—a generalization of QUBO—with GAS. The objective function and constraints need to be efficiently encoded, for which we use a *Quantum Dictionary* [20], a pattern for representing key-value pairs as entangled quantum registers, that turns out to be efficient for polynomial functions – in particular for quadratics representing QUBO problems. The approach relies on the addition of classical numbers to a quantum register in superposition, conditioned on the state of another quantum register. It is similar to the method used in *Quantum Fourier Transform* (QFT) adders [21]. Given a boolean polynomial, the coefficient of each monomial is added to the value register conditioned on the qubits in the key register corresponding to the variables present in the monomial.

We test our algorithm on the portfolio optimization problem [22, 23]. Multiple variants of this problem have been studied in the quantum optimization literature, ranging from convex continuous formulations [24] to QUBOs [13, 25, 26]. Here we investigate a QUBO formulation as well as a formulation with an inequality budget constraint, the latter not

being compatible with other approaches like quantum annealing, VQE, or QAOA—as those approaches can only handle linear equality constraints through quadratic penalty terms.

The remainder of this paper is organized as follows. Sec. 2 introduces GAS in general. Sec. 3 introduces QUBO problems, and shows how we can efficiently generate oracles to solve them using GAS algorithms, as well as how this approach extends to more general CPBO problems. In Sec. 4, we apply the developed technique to a concrete test case – portfolio optimization – and demonstrate it via simulation using Qiskit [27]. Sec. 5 concludes this paper and discusses possible directions of future research.

2 Grover Adaptive Search

Optimization problems are often solved by sequential approximation methods. In many cases, such methods are the only choice, but they may be computationally more efficient even when a solution to a problem can be expressed in a closed form. GAS works in a similar way, as it repeatedly uses Grover Search to randomly sample from all solutions better than the current one.

Grover Search is often described as a search algorithm, because it was initially formulated in the context of finding a single state of interest in a superposition of n -qubit quantum states. The algorithm has been generalized to the case of multiple states of interest, in which case it is better interpreted as a sampling algorithm. It amplifies the amplitudes of the states of interest within a larger search space, thus, increasing the probability of measuring one of the target states.

Grover Search – the core element of GAS – needs three ingredients:

1. A state preparation operator A to construct a superposition of all states in the search space. In this manuscript, A is implemented by Hadamard gates $H^{\otimes n}$, i.e. it constructs the equal superposition state:

$$H^{\otimes n} |0\rangle_n = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle_n. \quad (1)$$

2. An oracle operator O , that recognizes the states of interest and multiplies their amplitudes by -1. For instance, suppose $I \subset \{0, \dots, 2^n-1\}$ denotes the set of target states and $A = H^{\otimes n}$, then

$$OA|0\rangle_n = \frac{1}{\sqrt{2^n}} \sum_{i \notin I} |i\rangle_n - \frac{1}{\sqrt{2^n}} \sum_{i \in I} |i\rangle_n. \quad (2)$$

3. The Grover diffusion operator D , that multiplies the amplitude of the $|0\rangle_n$ state (or, equivalently, all states except $|0\rangle_n$) by -1.

Algorithm 1: Grover Adaptive Search

Input: $f : X \rightarrow \mathbb{R}$, $\lambda > 1$

```

1 Uniformly sample  $x_1 \in X$  and set  $y_1 = f(x_1)$ ;
2 Set  $k = 1$  and  $i = 1$ ;
3 repeat
4   Randomly select the rotation count  $r_i$  from
     the set  $\{0, 1, \dots, \lceil k-1 \rceil\}$ ;
5   Apply Grover Search with  $r_i$  iterations,
     using oracles  $A_{y_i}$  and  $O_{y_i}$ . We denote the
     outputs  $x$  and  $y$  respectively;
6   if  $y < y_i$  then
7      $x_{i+1} = x$ ,  $y_{i+1} = y$ , and  $k = 1$ 
8   else
9      $x_{i+1} = x_i$ ,  $y_{i+1} = y_i$ , and  $k = \lambda k$ 
10   $i = i + 1$ ;
11 until a termination condition is met;
```

The diffusion operator has the net effect of inverting all amplitudes in the quantum state about their mean. This causes all the amplitudes of the states of interest to be magnified, while the amplitudes of all other states are decreased. More precisely, applying the Grover operator $G = ADA^\dagger O$ the right number of times to state $A|0\rangle_n$ – i.e. evaluating $G^r A|0\rangle_n$ for an integer $r \geq 0$ – will maximally amplify the amplitudes of the states of interest. The optimal number of applications r depends on the number $N = 2^n$ of all states and the number s of states of interest, and is equal to $\lfloor \frac{\pi}{4} \sqrt{\frac{N}{s}} \rfloor$. This implies a probability of sampling a target state of at least 1/2, which corresponds to a quadratic speed-up compared to classical search. Since s is in general unknown, we can either use Quantum Counting algorithms [28–31] to find s , or apply a randomized strategy.

The latter is the essence of [32], where an algorithm for applying Grover Search for unknown s is presented. This was then used to create a minimum-finding algorithm [17], which we refer to as GAS. In the following we outline GAS, which is formally given in Alg. 1.

Consider a function $f : X \rightarrow \mathbb{R}$ for n binary variables, where for ease of presentation assume $X = \{0, 1\}^n$, for which we are interested in finding $\min_{x \in X} f(x)$. The main idea of GAS is to construct A_y and O_y for a given threshold y such that they flag all states $x \in X$ satisfying $f(x) < y$, such that we can use Grover Search to find a solution $\tilde{x}$ with a function value better than y . Then we set $y = f(\tilde{x})$ and repeat until some formal termination criteria is met, e.g. based on the number of iterations, time, or progress in y .

While implementations of GAS vary around the specific use case [19, 33], the general framework still loosely follows the steps described in [17]. In the following, we will show how operator A and oracle O can

be efficiently constructed for QUBO as well as CPBO problems.

3 QUBO and CPBO Oracles

A QUBO problem with n binary variables is specified by a quadratic operator represented by a matrix $Q \in \mathbb{R}^{n \times n}$, vector $b \in \mathbb{R}^n$, and constant $c \in \mathbb{R}$, defined as

$$\min_{x \in \{0,1\}^n} \left(\sum_{i,j=1}^n Q_{ij} x_i x_j + \sum_{i=1}^n b_i x_i + c \right), \quad (3)$$

or more compactly as $\min_{x \in \{0,1\}^n} (x^T Q x + b^T x + c)$, i.e. $f(x) = x^T Q x + b^T x + c$.

In the following, we show how to efficiently construct GAS oracles for QUBO problems. We will construct A_y such that it prepares a n -qubit input register to represent the equal superposition of all $|x\rangle_n$ and a m -qubit output register to (approximately) represent the corresponding $|f(x) - y\rangle_m$. Then, the oracle O_y should flag the states with a negative value in the output register. Note that in the implementation discussed, the oracle operator is actually independent of y , but this is not a requirement. For clarity, we will refer to the oracle as O when the oracle is independent of y . More formally, we show how to construct the oracles such that

$$A_y |0\rangle_n |0\rangle_m = \frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} |x\rangle_n |f(x) - y\rangle_m, \text{ and} \quad (4)$$

$$O |x\rangle_n |z\rangle_m = \text{sign}(z) |x\rangle_n |z\rangle_m \quad (5)$$

where $|x\rangle$ is the binary encoding of the integer x . Furthermore, we will show how the developed technique can be used to extend GAS to higher-degree polynomials of binary variables, as well as to constrained optimization.

3.1 Construction of operator A

To construct A , we will use a Quantum Dictionary, as introduced in [20], and summarize the construction in the following subsections.

3.1.1 Encoding a Single Integer Value

Given an m -qubit register and an angle $\theta \in [-\pi, \pi)$, we wish to prepare a quantum state whose state vector represents a "periodic signal" equivalent to a geometric sequence of length 2^m . This can be implemented using a unitary operator defined by Fig. 1.

The simplest implementation of $U_G(\theta)$ uses the phase gate $R(\theta)$ that, when applied to a qubit, rotates the phase of the amplitudes of the states having 1 in the position corresponding to that qubit. In Qiskit, this gate is the $U_1(\theta)$ operator [27]. The circuit for $U_G(\theta)$ is shown in Fig. 2, and consists of applying the

$$H^{\otimes m} |0\rangle_m \xrightarrow{U_G(\theta)} = \frac{1}{\sqrt{2^m}} \sum_{k=0}^{2^m-1} e^{ik\theta} |k\rangle_m$$

Figure 1: Definition of unitary operator $U_G(\theta)$, where $\theta \in [-\pi, \pi)$, applied to an m -qubit register in equal superposition. The result is a quantum state vector that represents a geometric sequence of length 2^m .

gate $R(2^i\theta)$ to the qubit $m-1-i$ in the m -qubit register prepared in the state of equal superposition in equation (1).

Figure 2: A circuit for unitary operator $U_G(\theta)$, where $\theta \in [-\pi, \pi)$, applied to an m -qubit register. R denotes the phase gate, which rotates the amplitudes of states having 1 in the position corresponding to the qubit it was applied to.

In [20] an alternative way to implement U_G was introduced by applying controlled R_y rotations to an ancillary register containing the encoding of an eigenstate of R_y , as explained in Appendix A.

Note that QFT applied to a register containing the binary encoding of a non-negative integer also creates a geometric sequence of amplitudes in that register. U_G can be seen as a shortcut for the QFT when the encoded numbers are known classically, as we avoid multi-qubit interactions.

Given an integer $-2^{m-1} \leq k < 2^{m-1}$, if we apply $U_G(\frac{2\pi}{2^m}k)$, followed by the inverse QFT to an m -qubit register prepared in the state of equal superposition in equation (1), we end up with $k \pmod{2^m}$ being encoded in the register, as shown in Fig. 3.

$$H^{\otimes m} |0\rangle_m \xrightarrow{U_G(\frac{2\pi}{2^m}k)} \xrightarrow{QFT^\dagger} = |k \pmod{2^m}\rangle_m$$

Figure 3: The geometric sequence encoding of an integer $-2^{m-1} \leq k < 2^{m-1}$, applied to a register of m qubits in equal superposition.

This representation is called the binary Two's Complement of k , which just adds 2^m to negative values k , similar to the way we can represent negative angles with their complement, e.g. equating $-\pi/4$ with $7\pi/4$. The reason this representation occurs naturally in this context is due to the fact that rotation composition behaves like modular arithmetic. The same method can be used to encode non-integers, as discussed in Appendix B.

Figure 4: An example of $C^{\{1,3\}}(U)$ with $n = 4$ input qubits and m output qubits.

3.1.2 Encoding a Superposition of Polynomial Values

Next we will discuss the application of $U_G(\theta)$ to a register $|z\rangle_m$ representing the values of a function, controlled on a register $|x\rangle_n$ representing the inputs of the same function. In general, the application of a unitary operator U to $|z\rangle_m$ controlled by a set $J \subseteq \{0, \dots, n-1\}$ of qubits of $|x\rangle_n$ can be expressed as

$$C^J(U) |x\rangle |z\rangle = |x\rangle U_{j \in J}^{x_j} |z\rangle, \quad (6)$$

and an example is shown in Fig. 4.

The general form of a polynomial of n variables is:

$$P(x) = \sum_{J \subseteq \{0, \dots, n-1\}} a_J \prod_{j \in J} x_j. \quad (7)$$

Each subset $J \subseteq \{0, \dots, n-1\}$ has a corresponding monomial $\prod_{j \in J} x_j$. The free term is represented by $a_\emptyset$.

3.1.3 Construction of Operator A

Now we are ready to define the operator A , as shown in Fig. 5. It consists of applying a controlled geometric sequence transformation $C^J(U_G(\frac{2\pi}{2^m} a_J))$ for each subset $J \subseteq \{0, \dots, n-1\}$ with a non-zero coefficient a_J , followed by a single application of the inverse QFT.

Figure 5: The circuit for state preparation operator A , applied to input register $|x\rangle_n$ and output register $|z\rangle_m$. Starting in a state of equal superposition, the operator employs several controlled applications of the unitary operator U_G , whose angle parameter corresponds to a non-zero coefficient a_J , where J is a subset of $\{0, \dots, n-1\}$. The single application of the inverse Quantum Fourier Transform ($QFT^\dagger$) at the end of the circuit decodes the periodic signal encoded by U_G , resulting in a superposition of key-value pairs.

Note that QUBO polynomials have only monomials of degree less than or equal to 2, i.e. $|J| \leq 2$, so we only need to control on single and pairs of qubits. An example of encoding a single monomial is shown in Fig. 6.

Figure 6: An example of encoding the monomial $2x_1x_3$, with input register $|x\rangle_4$ and output register $|z\rangle_3$. R denotes the phase gate, which rotates the amplitudes of states having 1 in the position corresponding to the qubit it was applied to.

Gate	Count
H	m
R	m
1-controlled R	$m \cdot n$
2-controlled R	$m \cdot n(n-1)/2$
Inverse QFT on m qubits	1

Table 1: Number of required gates to realize A in terms of number of input qubits n and output qubits m assuming a QUBO with dense Q and b and non-zero offset c . The depth scales as $1/m$ -times the number of gates since they can be mostly applied in parallel. Note that a QFT on m qubits can be implemented using $O(m \log(m))$ gates and thus will not dominate the overall circuit complexity [34]

The operator A prepares a state where the $|x\rangle_n$ register holds all 2^n inputs in equal superposition, entangled with the corresponding values $P(x)$ encoded in the $|z\rangle_m$ register:

$$A |0\rangle_n |0\rangle_m = \frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} |x\rangle_n |P(x)\rangle_m, \quad (8)$$

where we are assuming that an appropriate size m for the value register is known.

The desired A_y operator is obtained by adding the $-y$ constant to the free term of the original polynomial. This construction works for polynomials of arbitrary degree.

A detailed summary of the number of gates required to construct A for a QUBO is given in Table 1. In general, the number of gates scales as the number of monomials in the polynomial to be loaded. This matches the scaling of the description of the problem under the assumption that the input weights and output are represented using roughly the same number of bits, thus, the asymptotic scaling of our approach is optimal.

An alternative approach to construct A would be to use quantum arithmetic. However, even uncontrolled in-place addition of a classically given m -bit number to a m -bit quantum register requires $2m$ qubits and $2m-1$ Toffoli gates [35–37]. This then would have to be done $O(n^2)$ times in a 2-controlled way, which leads to significantly larger circuits and m more qubits than

Figure 7: A constrained optimization circuit, which encodes the functions corresponding to a given set of constraints into an m -qubit register, and flips a related indicator qubit for each condition that is satisfied. In this example, given a function $f : X \rightarrow \mathbb{R}$ of n binary variables, the global indicator qubit (shown at the bottom of the circuit) is set to $|1\rangle$ if and only if $f(x) + y$ —where y is the threshold parameter from GAS—and the cost of the input (denoted $C(x)$) are both less than 0. Note that the encoding method is the same as described in Sec. 3.1, meaning that the binary representation of the values is in Two’s Complement, and thus we only need to control on a single qubit (the most significant bit) to determine if a value is negative. There is a trade-off between reusing the m qubits and reversing the computation (shown here), or adding additional value qubits for each constraint, allowing one to skip the uncompute. In either case, we can replace the oracle that flags "good states" by a multi-controlled Z gate, with a number of controls equal to one more than the number of constraints. Here, we add one qubit and use a Toffoli gate to get back to the setting introduced before.

required by our approach, and also requires explicit encoding of negative integers.

3.2 Construction of oracle O

At each step of the algorithm, we are adding a constant to the polynomial, and searching for remaining negative values. This means the oracle just needs to recognize negative integers. Since values are represented in Two’s Complement, where the most significant (left-most) bit designates the sign of the number, a single qubit in the value register can be used to recognize negative integers. The typical oracle that multiplies target amplitudes by -1 can be applied. Note that the oracle stays unchanged between iterations, because we add a constant to the polynomial, which may lead to overflow in the value register. In order to avoid that, we may need to increase the number of qubits in the value register by 1.

Alternatively, threshold-based oracles (which are potentially more expensive) can be used, that will reduce the search space by filtering the numbers above a given threshold.

3.3 Constrained Optimization

It is common for optimization problems to impose additional constraints—e.g. the total number or cost of assets in a portfolio may be subjected to an upper bound. Such constraints translate into further reductions of the search space based on the key register. For example, the number of assets in a portfolio corresponds to the number of 1s in the binary representation of the input of the objective function, called the Hamming weight.

It is straightforward to use the GAS oracles introduced in Sec. 3.1 and 3.2 to take into account polynomial equality and inequality constraints on the key

register. Therefore, we can add additional registers to evaluate other polynomials. Whether an inequality constraint is satisfied or not can again be mapped to the sign qubit by applying an appropriate shift to the polynomial. Equality constraints are a bit more expensive, as they require the detection of a particular state, which essentially has the same complexity as the Grover diffusion operator D .

This leads to a set of qubits flagging target states: one qubit identifying the states that correspond to objective values below the current threshold, and one qubit for each constraint. Applying a logical AND-operation to all of them essentially acts as an intersection of the individually flagged states and allows to construct oracles for CPBO. An example is shown in Fig. 7, which we demonstrate in Sec. 4.1.

4 Test Cases

In the remainder of this paper we demonstrate the proposed technique on the portfolio optimization problem, and then show a simple example on quantum hardware.

4.1 Portfolio Optimization

Suppose an investment universe consisting of n assets, denoted by $i = 1, \dots, n$, their corresponding expected returns $\mu \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$. Furthermore, we consider a given risk factor $q \geq 0$, which determines the considered risk appetite. The resulting objective function is

$$\min_{x \in \{0,1\}^n} (qx^T \Sigma x - \mu^T x). \quad (9)$$

In other words, we want to minimize the weighted variance minus the expected portfolio return. Setting

$q = 0$ implies a risk neutral investor, while increasing q increases its risk averseness.

In the presented form, portfolio optimization is a QUBO problem. We can extend it by imposing a budget constraint of the form

$$\sum_{i=1}^n x_i = B, \quad (10)$$

where $B \in \{0, \dots, n\}$ denotes the number of assets to be chosen.

In general, equality constraints can be recast as penalty terms

$$\lambda \left(\sum_{i=1}^n x_i - B \right)^2, \quad (11)$$

and added to the objective (since we minimize), where $\lambda > 0$ is a large number to enforce the constraint to be satisfied. This results in a quadratic term that will again lead to a QUBO problem.

However, with the methodology introduced in Sec. 3.3, we can also model more complex constraints, e.g. a budget inequality constraint of the form

$$\sum_{i=1}^n c_i x_i \leq B, \quad (12)$$

where $c_i \in \mathbb{R}$ denote the asset prices, which does usually make more sense in practice than (10). In the following we consider two examples, one with an objection function that we want to minimize, and then the same problem with added constraints.

4.1.1 Finding a Minimum Value

Consider a portfolio of $n = 3$ assets, risk factor $q = 0.5$, and returns described by:

$$\mu = \begin{pmatrix} 1 \\ -2 \\ 3 \end{pmatrix} \quad \text{and} \quad \Sigma = \begin{pmatrix} 2 & 0 & -4 \\ 0 & 4 & -2 \\ -4 & -2 & 10 \end{pmatrix} \quad (13)$$

which leads to the formulation

$$\min_{x \in \{0,1\}^3} (-2x_1x_3 - x_2x_3 - 1x_1 + 2x_2 - 3x_3). \quad (14)$$

The objective function with added constant $-y$ (where y is the current threshold) has an associated A_y operator and the oracle O recognizes negative values, as introduced in Sec. 3. To perform the experiment we need 7 qubits split into two registers, $n = 3$ input qubits and 4 output qubits. While we only need 3 qubits in the output register, we add 1 extra to accommodate for the threshold shift. We set the initial threshold $y_1 = 0$, and stop searching if an improvement has not been found in three consecutive iterations of the algorithm.

For each iteration of GAS, we determine the number of applications of the Grover iterate as defined in Alg. 1. We apply A_y for the current threshold, and

Figure 8: The output probabilities of GAS for three iterations, searching for a minimum value of a function (Eq. 14), each with thresholds y and r applications of the Grover operator.

then apply the Grover iterate $G_y = A_y D A_y^\dagger O$ for the predetermined number of applications. If the measured value is less than y , we update the threshold. This process repeats until we have seen no improvement for three consecutive iterations.

Classically, we can determine the original minimum value by keeping track of the total shift, or by calculating the value of the objective function for the minimum key. The results of this simulated experiment are shown in Fig. 8.

4.1.2 Additional Constraints

We can impose a budget inequality constraint of the form discussed in Eq. 12 to the previous problem, where $B < 2$ and the price of each asset is 1. As shown in Sec. 3.3, we can implement this constraint by adding an additional register to the existing quantum circuit, and then encode the Hamming weight of the binary representation of the keys in that register. Note that we only need to control on the most significant qubit of the constraint register to determine if $B < 2$. Otherwise, the procedure for applying GAS is the same as in the original problem. The results of this simulated experiment are shown in Fig. 9.

4.2 Trials on Real Hardware

The experiments discussed in this section were run on the IBM ibmq_toronto device, with Quantum Volume 32 [38]. The configuration details for ibmq_toronto at the time of our experiments is given in Appendix C. Readout error-mitigation techniques were applied to the results of each circuit [27, 39].

Figure 9: The output probabilities of three iterations of GAS, with respective thresholds y and r applications of the Grover operator. We want to find a minimum value of Eq. 14, with the additional constraint that the binary representation of the corresponding key has a Hamming weight less than 2.

Let us consider a simple example of a polynomial minimization which can be run on current quantum hardware:

$$\min_{x \in \{0,1\}^2} (-2 + x_1 + x_2). \quad (15)$$

As in the previous subsection, we set the initial threshold $y_1 = 0$, and stop searching after three iterations with no improvement. Note that due to the probabilistic nature of the algorithm there is a non-zero probability that an invalid key-value pair will be measured. In addition, the noise inherent in the present era of quantum hardware further impacts the results and increases the probability of wrong results. We repeat the computation several times, and take the outcome with the maximum probability as the measured result. As long as the noise is not too strong, this still achieves a good key-value pair.

The results of the experiment are shown in Fig. 10. Note that in this example, the valid measurement outcomes are $0 \rightarrow -2$, $1 \rightarrow -1$, $2 \rightarrow -1$, and $3 \rightarrow 0$. In the first iteration the four outcomes are shown in an approximately-equal superposition, and sampled randomly (we do not apply the Grover iterate). We measure $1 \rightarrow -1$, and thus we update the threshold to $y_2 = -1$ and the number of iterations to $r = 1$. In the second iteration of Grover Adaptive Search the results of the amplification are shown, where $0 \rightarrow -2$ (the minimum) has the highest probability.

Figure 10: The output probabilities of GAS for two iterations run on real quantum hardware, with respective thresholds y and r applications of the Grover operator. The valid measurement outcomes are shown in blue, while the invalid measurement outcomes are shown in grey.

5 Conclusion

In this paper we introduced an efficient way to implement the oracles required for solving *Constrained Polynomial Binary Optimization* problems using *Grover Adaptive Search*. This problem class is very general and contains for instance QUBO problems. Our approach significantly reduces the number of gates required compared to standard quantum arithmetic approaches, i.e. it lowers the requirements to apply GAS on real quantum hardware for practically relevant problems. We demonstrated our algorithm on the portfolio optimization problem, i.e. a QUBO, where we could reliably find the optimal solution, and on real quantum hardware. Within this manuscript we focused mainly on problems with integer coefficients. The handling of non-integers is discussed in Appendix B.

Code Availability

All code associated with this publication can be found in the IBM Qiskit Optimization module, currently hosted at <https://github.com/Qiskit/qiskit-optimization>.

Acknowledgments

This material is for informational purposes only and is not the product of JPMorgan Chase & Co.'s Research Department. This material is not intended as research, a recommendation, advice, offer or solicitation for the purchase or sale of any financial product or service, and is not a research report and is not intended as such. This material is not intended to represent any position or opinion of JPMorgan Chase &

Figure 11: An example of an integer encoding (left), accompanied by two examples of a non-integer encoding (middle and right). Note that in the integer case only one outcome is possible, whereas the encoding of a non-integer leads to a superposition of approximations.

Co. JPMorgan Chase & Co. disclaims any responsibility or liability whatsoever for the quality, accuracy or completeness of the information herein, and for any reliance on, or use of this material in any way. ©2021 JPMorgan Chase & Co.

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. The current list of IBM trademarks is available at <https://www.ibm.com/legal/copytrade>.

A Quantum Dictionary

The *Quantum Dictionary* was introduced in [20] as a quantum computing pattern for encoding functions, in particular polynomials, into a quantum state using geometric sequences. The paper shows how quantum algorithms like search and counting applied to a quantum dictionary allow to solve combinatorial optimization and QUBO problems more efficiently than using classical methods.

Figure 12: Quantum Dictionary circuit.

Encoding a geometric sequence can be done using the phase gate as we explained earlier, but it is worth mentioning an alternative method described in [20], which uses the R_y family of gates.

$$|0\rangle \xrightarrow{R_x(\pi/2)} \xrightarrow{Z} \xrightarrow{X} = \xrightarrow{E(R_y)}$$

Figure 13: Eigenstate preparation for R_y .

The R_y gate is applied to an ancillary register containing one of its eigenstates prepared by an operator

$E(R_y)$, conditioned on both the key and value registers. The rotation angle θ is different for each application, representing a number that contributes to the values corresponding to keys that have the conditioned key as a subset. In particular, when encoding a polynomial, a rotation will be applied for each of its coefficients.

In [20] we used the R_y operator, with the eigenstate $1/\sqrt{2}(i|0\rangle + |1\rangle)$, independent of the rotation angle. This state can be prepared by the circuit in Fig. 13. The corresponding eigenvalue of $R_y(2\theta)$ is $e^{i\theta}$.

B Handling Non-Integers

When handling non-integers, we have two choices. The first is to approximate them with fractions with a common denominator and encode the numerator into quantum registers before performing computations, cf. Appendix B.1. The second, cf. Appendix B.2, is to phase-encode real numbers and let the inverse QFT convert the result of the computation into a superposition of approximations.

B.1 Approximating Real Coefficients by Fractions

If we relax the assumption that all coefficients are integers, we can approximate non-integers by dividing all values in μ and Σ by the largest (scaling the range of the coefficients to $[-1, 1)$), and approximating each value k by a fraction $\frac{k}{2^m}$ with $-2^{m-1} \leq k < 2^{m-1}$, where m is the number of value qubits. As an example, suppose $n = 3$, $m = 5$, $q = 0.5$, and

$$\mu = \begin{pmatrix} -3.77 \times 10^{-3} \\ 1.09 \times 10^{-3} \\ 2.41 \times 10^{-3} \end{pmatrix}. \quad (16)$$

Scaling the coefficients of μ leads to

$$\mu = \begin{pmatrix} -1.0 \\ 0.29 \\ 0.64 \end{pmatrix}, \quad (17)$$

and approximating them by fractions leads to

$$\mu = \begin{pmatrix} -16/32 \\ 5/32 \\ 10/32 \end{pmatrix}. \quad (18)$$

As the approximated function coefficients have a common denominator, we can ignore the denominator and treat the values as integers

$$\mu = \begin{pmatrix} -16 \\ 5 \\ 10 \end{pmatrix}, \quad (19)$$

which results in the optimization problem

$$\min_{x \in \{0,1\}^3} -(-16x_1 + 5x_2 + 10x_3). \quad (20)$$

B.2 Encoding Real Coefficients as Fejér Distributions

Recall the unitary operator $U_G(\theta)$ from Fig. 1, where $\theta \in [-\pi, \pi)$. As discussed in Sec. 3.1, the process for encoding an integer $-2^{m-1} \leq k < 2^{m-1}$ is to apply $U_G(\frac{2\pi}{2^m}k)$ to an m -qubit register in equal superposition, followed by a single application of the inverse Quantum Fourier Transform.

$$H^{\otimes m} |0\rangle_m \xrightarrow{U_G(\theta)} \xrightarrow{QFT^\dagger} = U_{\text{Fejér}}(\theta) |0\rangle_m$$

Figure 14: $U_G(\theta)$ followed by inverse QFT .

For the general case of a real number, shown in Fig. 14, where angle $\theta \in [-\pi, \pi)$. The application of the same sequence of gates from the integer case results in a quantum state whose state vector consists of the inner product between $G(\theta)$ and the Fourier bases $G(\frac{2\pi}{2^m}j)$, representing a similarity measure between θ and $\frac{2\pi}{2^m}j$, for $0 \leq j \leq 2^m - 1$, where $G(\theta)$ denotes the geometric sequence vector $(1, e^{i\theta}, \dots, e^{i(2^m-1)\theta})$. We will call the operator preparing this state $U_{\text{Fejér}}$ because the outcome probability distribution is the Fejér distribution [40]:

$$U_{\text{Fejér}}(\theta) |0\rangle_m = \frac{1}{\sqrt{2^m}} \sum_{j=0}^{2^m-1} \langle G(\theta), G(\frac{2\pi}{2^m}j) \rangle |j\rangle \quad (21)$$

If $\theta = \frac{2\pi}{2^m}a$, for a real number $-2^{m-1} \leq a < 2^{m-1}$, then $U_{\text{Fejér}}(\theta) |0\rangle_m$ prepares a state whose two most likely measurement outcomes are the closest two integers to a . The probability of measuring one of them is at least 81% [41]. Fig. 11 shows the probability distribution of the outcomes for multiple values of a where $m = 4$.

C Hardware Specifications

The error map for `ibmq_toronto` in Fig. 15 was generated on the day of experimentation using Qiskit’s visualization library [27].

Figure 15: The error map for `ibmq_toronto`.

References

- [1] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC ’96, pages 212–219, New York, NY, USA, 1996. ACM. ISBN 0-89791-785-5. DOI: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866).
- [2] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, Oct 1997. ISSN 1095-7111. DOI: [10.1137/s0097539795293172](https://doi.org/10.1137/s0097539795293172).
- [3] B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri, and et al. Towards quantum chemistry on a quantum computer. *Nature Chemistry*, 2(2):106–111, Jan 2010. ISSN 1755-4349. DOI: [10.1038/nchem.483](https://doi.org/10.1038/nchem.483).
- [4] Aram W. Harrow, Avinandan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15), Oct 2009. ISSN 1079-7114. DOI: [10.1103/physrevlett.103.150502](https://doi.org/10.1103/physrevlett.103.150502).
- [5] Carlos Bravo-Prieto, Ryan LaRose, M. Cerezo, Yigit Subasi, Lukasz Cincio, and Patrick J. Coles. Variational quantum linear solver, 2020. URL <https://arxiv.org/abs/1909.05820>.
- [6] Patrick Rebentrost, Brajesh Gupta, and Thomas R. Bromley. Quantum computational finance: Monte carlo pricing of financial derivatives. *Phys. Rev. A*, 98:022321, Aug 2018. DOI: [10.1103/PhysRevA.98.022321](https://doi.org/10.1103/PhysRevA.98.022321).
- [7] Stefan Woerner and Daniel J. Egger. Quantum risk analysis. *npj Quantum Information*, 5:15, 2019. DOI: [10.1038/s41534-019-0130-6](https://doi.org/10.1038/s41534-019-0130-6).
- [8] D. J. Egger, R. Garcia Gutierrez, J. Cahue Mestre, and S. Woerner. Credit risk analy-

- sis using quantum computers. *IEEE Transactions on Computers*, pages 1–1, 2020. DOI: [10.1109/TC.2020.3038063](https://doi.org/10.1109/TC.2020.3038063).
- [9] Nikitas Stamatopoulos, Daniel J. Egger, Yue Sun, Christa Zoufal, Raban Iten, Ning Shen, and Stefan Woerner. Option pricing using quantum computers. *Quantum*, 4:291, Jul 2020. ISSN 2521-327X. DOI: [10.22331/q-2020-07-06-291](https://doi.org/10.22331/q-2020-07-06-291).
 - [10] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A Quantum Approximate Optimization Algorithm, 2014. URL <https://arxiv.org/abs/1411.4028>.
 - [11] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man Hong Yung, Xiao Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5:4213, 2014. ISSN 20411723. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
 - [12] Giacomo Nannicini. Performance of hybrid quantum-classical variational heuristics for combinatorial optimization. *Physical Review E*, 99:013304, Jan 2019. DOI: [10.1103/PhysRevE.99.013304](https://doi.org/10.1103/PhysRevE.99.013304).
 - [13] Panagiotis Kl. Barkoutsos, Giacomo Nannicini, Anton Robert, Ivano Tavernelli, and Stefan Woerner. Improving variational quantum optimization using cvar. *Quantum*, 4:256, Apr 2020. ISSN 2521-327X. DOI: [10.22331/q-2020-04-20-256](https://doi.org/10.22331/q-2020-04-20-256).
 - [14] L. Braine, D. J. Egger, J. Glick, and S. Woerner. Quantum algorithms for mixed binary optimization applied to transaction settlement. *IEEE Transactions on Quantum Engineering*, 2: 1–8, 2021. DOI: [10.1109/TQE.2021.3063635](https://doi.org/10.1109/TQE.2021.3063635).
 - [15] Andrew M Childs, Edward Farhi, and John Preskill. Robustness of adiabatic quantum computation. *Physical Review A*, 65(1):012322, 2001. DOI: [10.1103/PhysRevA.65.012322](https://doi.org/10.1103/PhysRevA.65.012322).
 - [16] Mark W Johnson, Mohammad HS Amin, Suzanne Gildert, Trevor Lanting, Firas Hamze, Neil Dickson, Richard Harris, Andrew J Berkley, Jan Johansson, Paul Bunyk, et al. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, 2011. DOI: [10.1038/nature10012](https://doi.org/10.1038/nature10012).
 - [17] Christoph Dürr and Peter Høyer. A quantum algorithm for finding the minimum, 1996. URL <https://arxiv.org/abs/quant-ph/9607014>.
 - [18] D. Bulger, W. P. Baritompa, and G. R. Wood. Implementing pure adaptive search with grover’s quantum algorithm. *Journal of Optimization Theory and Applications*, 116(3): 517–529, Mar 2003. ISSN 1573-2878. DOI: [10.1023/A:1023061218864](https://doi.org/10.1023/A:1023061218864).
 - [19] W. P. Baritompa, D. W. Bulger, and G. R. Wood. Grover’s quantum algorithm applied to global optimization. *SIAM J. on Optimization*, 15(4):1170–1184, April 2005. ISSN 1052-6234. DOI: [10.1137/040605072](https://doi.org/10.1137/040605072).
 - [20] Austin Gilliam, Charlene Venci, Sreraman Muralidharan, Vitaliy Dorum, Eric May, Rajesh Narasimhan, and Constantin Goniculea. Foundational patterns for efficient quantum computing, 2019. URL <https://arxiv.org/abs/1907.11513>.
 - [21] Thomas G. Draper. Addition on a quantum computer, 2000. URL <https://arxiv.org/abs/quant-ph/0008033>.
 - [22] Román Orús, Samuel Muel, and Enrique Lizaso. Quantum computing for finance: Overview and prospects. *Reviews in Physics*, 4:100028, Nov 2019. ISSN 2405-4283. DOI: [10.1016/j.revip.2019.100028](https://doi.org/10.1016/j.revip.2019.100028).
 - [23] Daniel J. Egger, Claudio Gambella, Jakub Marecek, Scott McFaddin, Martin Mevissen, Rudy Raymond, Andrea Simonetto, Stefan Woerner, and Elena Yndurain. Quantum computing for finance: State-of-the-art and future prospects. *IEEE Transactions on Quantum Engineering*, 1:1–24, 2020. ISSN 2689-1808. DOI: [10.1109/tqe.2020.3030314](https://doi.org/10.1109/tqe.2020.3030314).
 - [24] Patrick Rebentrost and Seth Lloyd. Quantum computational finance: quantum algorithm for portfolio optimization, 2018. URL <https://arxiv.org/abs/1811.03975>.
 - [25] Davide Venturelli and Alexei Kondratyev. Reverse Quantum Annealing Approach to Portfolio Optimization Problems. *Quantum Machine Intelligence*, 1, 2019. DOI: [10.1007/s42484-019-00001-w](https://doi.org/10.1007/s42484-019-00001-w).
 - [26] Daniel J. Egger, Jakub Marecek, and Stefan Woerner. Warm-starting quantum optimization, 2020. URL <https://arxiv.org/abs/2009.10095>.
 - [27] Héctor Abraham et. al. Qiskit: An open-source framework for quantum computing. 2019. DOI: [10.5281/zenodo.2562110](https://doi.org/10.5281/zenodo.2562110).
 - [28] Michele Mosca. Quantum searching, counting and amplitude amplification by eigenvector analysis. In *In Proceedings of Randomized Algorithms, Workshop of Mathematical Foundations of Computer Science*, pages 90–100, 1998. URL <http://141.89.225.3/resources/pdf/moscafi.pdf>.
 - [29] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum Amplitude Amplification and Estimation. *Contemporary Mathematics*, 305, 2002. DOI: [10.1090/conm/305](https://doi.org/10.1090/conm/305).
 - [30] Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tomoki Tanaka, Tamiya Onodera, and Naoki Yamamoto. Amplitude estimation without phase estimation. *Quantum Information Processing*, 19(2), Jan 2020. ISSN 1573-1332. DOI: [10.1007/s11128-019-2565-2](https://doi.org/10.1007/s11128-019-2565-2).

- [31] Scott Aaronson and Patrick Rall. Quantum approximate counting, simplified. *Symposium on Simplicity in Algorithms*, page 24–32, Jan 2020. DOI: [10.1137/1.9781611976014.5](https://doi.org/10.1137/1.9781611976014.5).
- [32] Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. Tight bounds on quantum searching. *Fortschritte der Physik*, 46(4-5):493–505, Jun 1998. ISSN 1521-3978. DOI: [10.1002/\(SICI\)1521-3978\(199806\)46:4/5<493::AID-PROP493>3.0.CO;2-P](https://doi.org/10.1002/(SICI)1521-3978(199806)46:4/5<493::AID-PROP493>3.0.CO;2-P).
- [33] Benjamín Barán and Marcos Villagra. *Multi-objective Optimization Grover Adaptive Search*, pages 191–211. Springer International Publishing, Cham, 2019. ISBN 978-3-319-99648-6. DOI: [10.1007/978-3-319-99648-6_11](https://doi.org/10.1007/978-3-319-99648-6_11).
- [34] Yunseong Nam, Yuan Su, and Dmitri Maslov. Approximate quantum fourier transform with $\mathcal{O}(n \log(n))$ T gates. *npj Quantum Information*, 6(1), Mar 2020. ISSN 2056-6387. DOI: [10.1038/s41534-020-0257-5](https://doi.org/10.1038/s41534-020-0257-5).
- [35] Steven A. Cuccaro, Thomas G. Draper, Samuel A. Kutin, and David Petrie Moulton. A new quantum ripple-carry addition circuit, 2004. URL <https://arxiv.org/abs/quant-ph/0410184>.
- [36] Thomas G. Draper, Samuel A. Kutin, Eric M. Rains, and Krysta M. Svore. A logarithmic-depth quantum carry-lookahead adder. *Quantum Info. Comput.*, 6(4):351–369, July 2006. ISSN 1533-7146. DOI: [10.26421/QIC6.4-5-4](https://doi.org/10.26421/QIC6.4-5-4).
- [37] Craig Gidney. Halving the cost of quantum addition. *Quantum*, 2:74, Jun 2018. ISSN 2521-327X. DOI: [10.22331/q-2018-06-18-74](https://doi.org/10.22331/q-2018-06-18-74).
- [38] Andrew W. Cross, Lev S. Bishop, Sarah Sheldon, Paul D. Nation, and Jay M. Gambetta. Validating quantum computers using randomized model circuits. *Physical Review A*, 100(3), Sep 2019. ISSN 2469-9934. DOI: [10.1103/physreva.100.032328](https://doi.org/10.1103/physreva.100.032328).
- [39] A. Dewes, F. R. Ong, V. Schmitt, R. Lauro, N. Boulant, P. Bertet, D. Vion, and D. Esteve. Characterization of a two-transmon processor with individual single-shot qubit readout. *Phys. Rev. Lett.*, 108:057002, Feb 2012. DOI: [10.1103/PhysRevLett.108.057002](https://doi.org/10.1103/PhysRevLett.108.057002).
- [40] Svante Janson. Moments of gamma type and the brownian supremum process area. *Probab. Surveys*, 7:1–52, 2010. DOI: [10.1214/10-PS160](https://doi.org/10.1214/10-PS160).
- [41] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011. ISBN 1107002176, 9781107002173. DOI: [10.1017/CBO9780511976667](https://doi.org/10.1017/CBO9780511976667).